arXiv:1609.00161v2 [cs.DS] 2 Sep 2016

Parallel Clustering of Graphs for Anonymization and Recommender Systems

Frédéric Prost
ENS de Lyon, INRIA, UCB Lyon 1, Laboratoire LIP
46 Alle d'Italie 69364 LYON Cedex 07, FRANCE
frederic.prost@ens-lyon.fr

Jisang Yoon
ENS de Lyon, INRIA, UCB Lyon 1, Laboratoire LIP
46 Alle d'Italie 69364 LYON Cedex 07, FRANCE
Jisang.Yoon@ens-lyon.fr

September 5, 2016

Abstract

Graph clustering is widely used in many data analysis applications. In this paper we propose several parallel graph clustering algorithms based on Monte Carlo simulations and expectation maximization in the context of stochastic block models. We apply those algorithms to the specific problems of recommender systems and social network anonymization. We compare the experimental results to previous propositions.

1 Introduction

An important way to discover structural properties within data is to classify them [11], that is to regroup *similar* elements into classes, called clusters. In this paper we focus on graph data: for this special case the edges of the graph represent the correlations between nodes, hence it is the edge topology that is used to define the clusters of the graph. There are many ways to have a meaningfull definition of what means to be *similar* (see [14]).

The stochastic block model, see [10], is widely used for graph clustering. The basic idea is that if some nodes behave *similarly* in network, it can be assumed that those nodes behave the same *probabilistically* and form a *cluster* (or a *block*).

Example 1.1 *The simplest stochastic block model is depicted in* ***Figure 1****.*

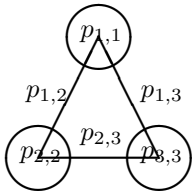

Figure 1: 3-clustering random graph model

The intuitive meaning of such a model is that $p_{i,j}$ denotes the probability of having an edge between a vertex in cluster i and a vertex in cluster j.

*There is a large class of graphs that fits this model. Let us assume that there are k-partition of vertices, $(V_i)_{i\in[k]}$ (**Notation**: $[n] := \{0, 1, \cdots, n\}$ for $n \in \mathbb{N}$). Any graph G such that for all $u \in V_i, v \in V_j (0 \leq i \leq j < k, u \neq v)$, there is an edge between u, v with probability $0 \leq p_{i,j} \leq 1$ fits the model.*

If we consider the problem the other way around, we have a concrete graph G for which we look for a SBM of it that is optimal, that is the optimal clustering $Z : V(G) \to [k]$, and parameters $p_{i,j}$'s.

The goal in **Example 1.1** is to find the most *probable* clustering for a given observed network. This kind of approach is known as the *Maximum Likelihood problem*. Here, the parameter is the clustering Z (discrete value), and the edge probability $p'_{i,j}s$ (continuous value).

In the most general case, finding the global optimum is a NP-hard problem, hence untractable (especially when dealing with very large data set coming from social networks for instance). However, there exists a general greedy algorithm (finding the *local* optimum, iteratively) which is called *Expectation-Maximization (EM)* algorithm [5] that works well in many problems [16]. Unfortunalely, there are some difficulties to apply this algorithm to the problem of graph clustering. In this paper we suggest a new algorithm to deal with this particular case. To the best of the authors knowledge, there was no similar algorithm presented in the literature.

In order to assess the interest of our algorithms, we show that they work well in practice: both for the quality of the answers and the time cost. We have considered two case studies: recommender system, and social network anonymization. For the recommender system case, the idea is to make clusters of similiar users with relation to products. For the social network anonymization case, the idea is to build a SBM of the graph by considering clusters of size k. Then an anonymized graph is created by re-expanding the clusters along the probabilities of the SBM: we expect to generate a graph with similar statistical properties of the original one in which it is not possible to reidentify nodes with an accuracy less than k.

We begin in section 2 by recalling basic definitions and results, and also defining our notations, on the Expectation-Maximization problems. In section 3 we define algorithms for the problem of graph clustering. In section 4, we define more precisely the two domains over which we have tested our algorithms: recommender systems and social network anonymization. In section 5 we explore the experimental behavior of our algorithms. We discuss related works in section 6, and finally we conclude in section 7.

2 Expectation-Maximization problems

Expectation-Maximization (EM) algorithm [6] denotes a large class of algorithms to tackle the problem of computing the maximum likelihood estimates from incomplete data. In this paper we apply this approach to the problem of graph clustering. We start by precisely defining the problem setting in its most generic way as well as our notations.

Definition 2.1 (Problem setting) *Let X be the observed data, Z be the unobserved latent data (in the case of discrete values we write $\mathcal{Z}$ the set of possible values of Z), and g be the probability mass function (pmf) of Z. Thus we derive $g(Z) = \int f_\theta(X, Z)$, where θ is a parameter of the distribution and $f_\theta(X, Z)$ is the probability density function (pdf) or pmf of X, Z.*

The goal is to maximize the log-likelihood function, that is to find

$$\hat{\theta} := \underset{\theta}{argmax}\ f_\theta(X) \tag{1}$$

where $f_\theta(X)$ is defined as

$$f_\theta(X) := \sum_{Z \in \mathcal{Z}} f_\theta(X, Z) \tag{2}$$

For computational convenience, we usually take the log to the likelihood function f_θ, and since such value is negative, we define the *entropy* as follows:

$$S(\theta) = -log \circ f_\theta \tag{3}$$

The entropy is used since we we may not be able to express f_θ as a float value for very large values of f_θ, e.g. $f_\theta \approx 10^{10^6}$.

Thus, from this point, our goal is to minimize entropy, instead of maximizing the log-likelihood function. The generic EM algorithm, for which local convergence is proved in [17], can be written as follows:

Algorithm 1: Generic EM

Input : observed data X, unobserved data Z, pmf $f_\theta(X, Z)$

Output: optimal $\hat{\theta}$

Pick random θ

repeat

  $g(Z) \leftarrow g(Z|X,\theta) := \frac{f_\theta(X,Z)}{\sum_{Z'} f_\theta(X,Z')}$ (E-step)

  $\theta \leftarrow \underset{\theta}{argmax}\ \mathbb{E}_g[log \circ f_\theta(X,Z)]$ (M-step)

until $S(\theta)$ *converges*;

Example 2.1 (k-means) *Regarding the way we consider the unobserved latent variable Z there are two types of EM: soft EM and hard EM. If we consider Z as a* variable*, we are in the case of a soft EM problem. If we consider Z as* parameter*, which we want to figure out, we are in the case of a hard EM problem. In hard EM, we assume that Z has a deterministic value.*

*The generic algorithm of a hard EM is represented as **Algorithm 2**.*

Algorithm 2: Generic Hard EM

Input : observed data X, unobserved latent data Z, pdf or pmf $f_{\theta,Z}(X)$

Output: optimal $\hat{\theta}$,$\hat{Z}$

Take random θ

repeat

  $Z \leftarrow \underset{Z'}{argmin}\ S(\theta, Z')$(E-step)

  $\theta \leftarrow \underset{\theta'}{argmin}\ S(\theta', Z)$(M-step)

until $S(Z)$ *converges*;

k-mean clustering is a particular case of hard EM. For this problem $X \subseteq \mathbb{R}^d$ is an unclassifed data set, and $Z : X \to [k]$ is a function which corresponds to a classification. The parameter θ is a k-tuple $(\mu_i \in \mathbb{R}^d)_{i\in[k]}$. The probability function $f_{\theta,Z}$ can be defined as follows:

$$f_{\theta,Z}(X) := \prod_{x\in X} (norm_dist(\mu_{Z(x)}, I_d, x)) \quad (4)$$

*where $norm_dist(\mu, \Sigma, x)$ is defined as the pdf of $\mathcal{N}(\mu, \Sigma)$. Thus in line 3 of **Algorithm 2**, for each $x \in X$, $Z(x)$ is assigned to $\underset{i\in[k]}{argmin}|x - \mu_i|^2$, and in line 4, μ_i is computed as the euclidean center of cluster $Z^{-1}(\{i\}) \subseteq X$. Simply sayed, k-mean clustering is a hard EM for a Gaussian mixture model with fixed covariance.*

3 Graph Clustering

3.1 Soft graph clustering (soft SBM)

We work on the problem of soft clustering for various kinds of graphs including simple graphs, weighted graphs *etc.* Each edge is assumed to have some configuration $r \in R$ (R is the set of configurations, *e.g.* existence, weight, rating, *etc*). Let us precisely define the problem setting of soft SBM.

Definition 3.1 (Soft SBM) *let X be the observed data (like the existence, the weight, the rating,* etc.*) of a given set of edges of a graph $G = (V, E)$. Let $h = (h_{u,i})_{u\in V(G), i\in[k]}$, where $\sum_{i\in[k]} h_{u,i} = 1, \forall u \in V(G)$, and $\theta = (\theta_{i,j,r})_{i,j\in[k],r\in R}$, where $\sum_{r\in R} \theta_{i,j,r} = 1, \forall i, j \in [k]$. Then the goal function $S(\theta, h)$ is defined as*

$$S(\theta, h) := -log \circ \prod_{(u,v)\in X} \Big(\sum_{i,j\in[k]} h_{u,i} \cdot h_{v,j} \cdot \theta_{i,j,r(u,v)} \Big) \quad (5)$$

where $r(u, v)$:=config of (u, v)(which is observed). The aim is to minimize $S(\theta, h)$.

Remark 1 *(θ, h) corresponds to θ in **Definition 2.1**. In addition, Z is defined as a function from X to $[k] \times [k]$, and $f_{\theta,h}(X, Z)$ is defined as:*

$$f_{\theta,h}(X, Z) = \prod_{(u,v)\in X} (h_{u,i} \cdot h_{v,j} \cdot \theta_{i,j,r(u,v)}) \quad (6)$$

where $Z(u, v) = (i, j) \in [k]^2$ for each $(u, v) \in X$.

[7] previously suggested an algorithm (**MMSBM**, Mixed Membership SBM) based on EM. Here, we propose a *randomized* algorithm (**MCMMSBM**,

Monte-Carlo MMSBM), for improving the complexity of **MMSBM**. The algorithm is based on a Monte-Carlo simulation and **MMSBM**. We discuss more precisely the relation between **MCMMSBM** and **MMSBM** in the related works section (see section 6).

Algorithm 3: MCMMSBM

Input : observed edge data X, $V(G)$:set of vertices, k: number of clusters, s: sample size

Output: optimal $\hat{\theta}, \hat{h}$

Take random θ, h
repeat
  for $(u,v) \in X, i \in [k]$ **do in parallel**
    $x_{u,v}(i), x_{v,u}(i) \leftarrow 0$
  end
  for $i,j \in [k], r \in R$ **do in parallel**
    $\eta_{i,j,r} \leftarrow 0$
  end
  for $(u,v) \in X$ **do in parallel**
    $isample \leftarrow discrete(h_u, s)$,
    $jsample \leftarrow discrete(h_v, s)$
    $x \leftarrow 0$
    $r \leftarrow r(u,v)$
    for $s' \in [s]$ **do**
      $i \leftarrow isample[s'], j \leftarrow jsample[s']$
      $x, x_{u,v}(i), x_{v,u}(j) += \theta_{i,j,r}$
    end
    for $s' \in [s]$ **do**
      $i \leftarrow isample[s'], j \leftarrow jsample[s']$
      $x_{u,v}(i), x_{v,u}(j) /= x$
      $\eta_{i,j,r} += \theta_{i,j,r}/x$
    end
  end
  for $u \in V(G), i \in [k]$ **do in parallel**
    $h_{u,i} \leftarrow \frac{1}{|\partial u|} \sum_{v \in \partial u} x_{u,v}(i)$
  end
  for $i,j \in [k], r \in R$ **do in parallel**
    $\theta_{i,j,r} \leftarrow \frac{\eta_{i,j,r}}{\sum_{r' \in I} \eta_{i,j,r'}}$
  end
until $S(\theta, h)$ *converges*;

∂u is defined as $\{v \in V(G) | (u,v) \in X\}$. The time complexity of **MCMMSBM** is $O(|X| \cdot (k + s \cdot log(k)))$. The sample size s is chosen by trade-off between cost and accuracy. Note that the time cost to generate s samples from an arbitrary discrete distribution with size k is $O(k + s \cdot log(k))$. It can be implemented by binary search in cumulative probabilities.

Experimental results (see **Section 5**) show that **MCMMSBM** and **MMSBM** produce similar results in terms of quality, but that **MCMMSBM** requires less resources than **MMSBM**.

3.2 Hard graph clustering

3.2.1 Hard classification problems and hard EM

Let us start by defining the problem setting of hard classification as follows:

Definition 3.2 *X: observed data, A: unclassified data, Z: unobserved latent classification from A to $[k]$, θ: parameter distribution, $f_{\theta,Z}(X)$: pmf of X. The goal is to minimize*

$$S(\theta, Z) := -log \circ f_{\theta,Z}(X) \tag{7}$$

For given Z, computing optimal $\hat{\theta} := \underset{\theta'}{argmin}\ S(\theta', Z)$ is usually not very expensive. Indeed, we can differentiate the entropy function for θ, since θ has continuous value while Z doesn't. Thus, in hard classification problems, we only consider Z as a parameter, and we can express entropy as

$$S(Z) := \min_{\theta'} S(\theta', Z) \tag{8}$$

The k-means algorithm (Gaussian mixture with fixed covariance) is also an algorithm for hard classification problems (**Example 2.1**). In the most general case (see the line 3 of **Algorithm 2**), the problem to solve boils down to the computation of such an *argmin* in general classification problems. The number of possible Z is equal to $|\mathcal{Z}| = k^{|A|}$. Thus, it is untractable to compare all the $k^{|A|}$ cases.

The Gaussian mixture model is just a special case in which we can easily compute *argmin* in line 3. Indeed, one can just compute argmin $\underset{Z(a)}{argmin} f_{\theta,X}(Z)$ for each $a \in A$.

However, computing the clustering Z for a given parameter θ is not easy. Indeed, deciding the best $Z(v)$ for each $v \in V(G)$ *depends* on the values of Z for others $v \in V(G)$. Thus, we can't decide $Z(v)$, separately as it is the case in a Gaussian mixture model. Hence, we proposed a new classification algorithm, *Generalized k-means*, which is inspired from the original k-means algorithm.

Algorithm 4: Generalized k-means

Input : observed data X, A : set of unclassified data, k : number of clusters, $\alpha \in (0, 1]$: iteration scale

Output: Z which minimizes $S(Z)$

1 Take random Z
2 **repeat**
3 $A' \leftarrow sample(A, \alpha|A|)$(uniform sample)
4 **for** $a \in A'$ **do in parallel**
5 $i \leftarrow \underset{i' \in [k]}{argmin}\ S(succ(Z, a, i'))$
6 Plan to reassign $Z(a) \leftarrow i$
7 **end**
8 Do plan, compute some information with X, Z(*e.g.* $\hat{\theta}$)
9 **until** $S(Z)$ *converges*;

Let $succ(Z, a, i) : A \to [k]$ be defined as a new classification satisfying:

$$succ(Z, a, i')(a') := \begin{cases} Z(a') & \text{if } a' \neq a \\ i' & \text{else} \end{cases} \tag{9}$$

Definition 3.3 *If no single movement between two clusters can improve the entropy of classification, then we call the such a classification a* locally optimal classification.

Generalized k-means algorithm achieves locally optimal classification, as same as classical k-means algorithm.

To compute $S(Z)$, we have to compute $\hat{\theta}(Z)$, and we can approximate

$$\begin{aligned} S(succ(Z, a, i')) &:= -log(f_{X,succ(Z,a,i')}(\hat{\theta}(succ(Z, a, i')))) \\ &\approx -log(f_{X,succ(Z,a,i')}(\hat{\theta}(Z))) \end{aligned} \tag{10}$$

if we assume that $|A| \gg 1$, and MLE of $f_{X,Z}(\theta)$ are *consistent*, this approximation works very well.

The third line of **Algorithm 4** is the key part of the algorithm. If α is set to 1.0, then this algorithm may not work for the hard graph clustering problem. Indeed, in this case it is not possible to assume that the proportion of movement during the parallel-loop is *negligible*.

Remark 2 *If $S(Z)$ starts to decrease in early iterations, it may continue to decrease (hence converge), because as $S(Z)$ decreases, most of $Z(a)$ for $a \in A'$ may keep its value, so that $\underset{i' \in [k]}{argmin} S(succ(Z, a, i'))$ becomes more accurate.*

3.2.2 Hard clustering of simple graph

Let us consider a real network as an observed data from a random graph model with clustering. Actually, such an approach constitutes a random graph model for *simple graphs*. It is possible to extend this approach to other kinds of graphs (*e.g.* directed graphs, weighted graphs, *etc.*).

Now, let us apply Generalized k-means on a hard simple graph clustering problem.

First, let us compute $S(Z)$.

$$\begin{aligned} S(Z) &= -\sum_{0 \leq i \leq j < k} (d_{i,j} log(\hat{p}_{i,j}) + d'_{i,j} log(1 - \hat{p}_{i,j})) \\ &= \sum_{0 \leq i \leq j < k} ((d_{i,j} + d'_{i,j}) log(d_{i,j} + d'_{i,j}) \\ &\quad - d_{i,j} log(d_{i,j}) - d'_{i,j} log(d'_{i,j})) \\ &\quad \text{(because MLE of Bernoulli distribution)} \\ &= \sum_{0 \leq i \leq j < k} f(d_{i,j}, d'_{i,j}) \end{aligned} \tag{11}$$

where

$V_i := Z^{-1}(\{i\})$

$$D_{i,j} := \begin{cases} |V_i|(|V_i|-1)/2 & \text{if } i = j \\ |V_i||V_j| & \text{else} \end{cases}$$

$d_{i,j} :=$ (# edges between $V_i \& V_j$)

$d'_{i,j} := D_{i,j} - d_{i,j}$

$$f(x,y) := \begin{cases} (x+y)log(x+y) - x \cdot log(x) - y \cdot log(y) & \text{if } x, y \geq 0 \\ 0 & \text{else} \end{cases} \quad (12)$$

Now, let us consider **Algorithm 5**.

Algorithm 5: Simple graph clustering via Generalized k-means

Input : G: Observed simple graph data, $Z : V(G) \to \{1, 2, \cdots, k\}$: unobserved latent classification data, $\alpha \in (0, 1]$: iteration scale

Output: Z which minimizes $S(Z)$

1 Take random Z
2 **repeat**
3 $V' \leftarrow sample(V(G), \alpha|V(G)|)$(uniform sample)
4 **for** $v \in V'$ **do in parallel**
5 $i \leftarrow \underset{i' \in [k]}{argmin}\ S(succ(Z, v, i'))$(**Algorithm 7**)
6 Plan to reassign $Z(v) \leftarrow i$
7 **end**
8 Do plan, compute $d_{i,j}, d'_{i,j}, x_{v,j}, A_i, a_{i,j}, B_i, b_{i,j}$ (**Algorithm 6**) for $0 \leq i \leq j < k, v \in V(G)$
9 **until** $S(Z)$ *converges*;

A_i represents the entropy increment corresponding to the merging of an isolated vertex to the cluster V_i. B_i represents the entropy decrement corresponding to the split of an isolated vertex from the cluster V_i. The reason for the computations of A_i, B_i is that the real network is seen as a sparse network. Thus, we can consider any arbitary vertex $v \in V'$ in **Algorithm 5** as an *almost* isolated vertex. If there are some edges from v to V_j, one just has to modify A_i. Then, the time complexity of line 5 in **Algorithm 5** is improved from $O(k^2)$ to $O(min(d \cdot k, k^2))$, where d is the average degree($= \frac{2m}{n}$). If $O(m) \approx O(n)$, then d

Algorithm 6: substep of **Algorithm 5**

1 Compute $D_{i,j}, d_{i,j}, d'_{i,j}, V_i$ as in Eq (12)
2 $x_{v,j} \leftarrow$ (# edges from v to V_j)
3 **for** $i \in [k]$ **do in parallel**
4 $A_i \leftarrow 0$
5 $B_i \leftarrow 0$
6 **for** $1 \leq j \leq k$ **do**
7 $a_{i,j} \leftarrow f(d_{i,j}, d'_{i,j} + |V_j|) - f(d_{i,j}, d'_{i,j})$
8 $b_{i,j} \leftarrow f(d_{i,j}, d'_{i,j}) - f(d_{i,j}, d'_{i,j} - |V_j|)$
9 $A_i \leftarrow A_i + a_{i,j}$
10 $B_i \leftarrow B_i + b_{i,j}$
11 **end**
12 **end**

Algorithm 7: substep of **Algorithm 5**

1 $temp \leftarrow B_y$
2 **for** j *s.t.* $x_{v,j} \neq 0$ **do**
3 $temp \leftarrow temp - b_{y,j}$
4 $y_{v,j} \leftarrow |V_j| - x_{v,j}$
5 $temp \leftarrow temp + f(d_{y,j}, d'_{y,j}) - f(d_{y,j} - x_{v,j}, d_{y,j} - y_{v,j})$
6 **end**
7 $min \leftarrow temp$
8 $argmin \leftarrow y$
9 **for** $y \in [k]$ **do**
10 $temp \leftarrow A_y$
11 **for** j *s.t.* $x_{v,j} \neq 0$ **do**
12 $temp \leftarrow temp - a_{y,j}$
13 $y_{v,j} \leftarrow |V_j| - x_{v,j}$
14 $temp \leftarrow temp + f(d_{y,j} + x_{v,j}, d_{y,j} + y_{v,j}) - f(d_{y,j}, d'_{y,j})$
15 **end**
16 **if** $temp < min$ **then**
17 $min \leftarrow temp$
18 $argmin \leftarrow y$
19 **end**
20 **end**

is a constant, thus the complexity for line 5 is $O(k)$.

Now, let's compute the global complexity of **Algorithm 5**. $\alpha \cdot s$ can be considered as the number of repeat-loop iterations, because all the vertices in $V(G)$ have to be correctly assigned to the clusters. Let m, n be the respecrtive numbers of edges and vertices in G. The for-loop in line 4 may take $O(min(m \cdot k, n \cdot k^2))$ by **Algorithm 5**, line 8 take $O(m + k^2)$, but by considering $m \gg k^2$, we have $O(m + k^2) \approx O(m)$. In conclusion, the complexity of **Algorithm 5** is $O(\frac{min(m \cdot k, n \cdot k^2)}{P})$, where P is the number of processors.

4 Applications

4.1 Recommender System

[3] One of the basic approach to recommender system is to consider the user-product relationship as a bipartite graph (see **Figure 2**). Users may rate each product they purchase or press the *like* button on some products. The former case can be considered as a weighted graph, and the latter case as a simple graph. The observed data is not the full graph, because each user might not experience/purchase all the products. Usually only a *partial* observation of the edges of the graph is considered. The goal is to anticipate the weight or the existence of hidden (unobserved) edges.

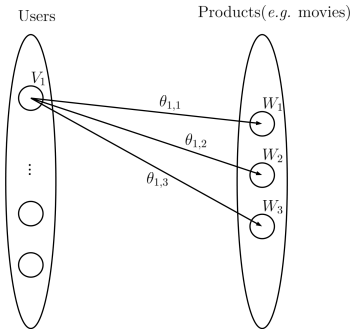

Figure 2: SBM for recommender system

In order to anticipate the weight or the existence of the edges, we look at this problem through the SBM point of view. Let us consider that there are groups of *similar users or similar products*. The existence or the weight of edges between clusters V_i, W_j follow some random distribution model (*e.g.* Bernoulli, binomial, *etc.*, see **Figure 2**). Using **MCMMSBM**, **MMSBM**, and Generalized k-means, the clustering can be optimized, and used to anticipate the weight or the existence of the hidden edges.

4.2 Social Network Anonymization

The information in social networks becomes an important data source, and sometimes it is necessary or beneficial to release such data to the public. Many real-world social networks contain sensitive information and serious privacy concerns on graph data have been raised. The famous result of Narayan and Shmatikov [13] has shown that naïve anonymization does not work: it is in practise very easy to re-identify elements of a trivially anonymized (ie replacing identifying informations such as names, social security numbers etc. with random numbers) social networks. Later works [2] pushed further the study of attacks on anonymized social networks.

The goal of social network anonymization is to produce a graph in such a way that some statistical functions produce the same result on the original graph and on the transformed graph, while other functions (namely reidentification) should not produce the same result. There are two main ways to work on the anonymization:

1. Clustering: one tries to group together edges and nodes so that when the the cluster regroups k nodes then there is no way to distinguish an individual node among them.

2. k-anonymity: one tries to modify the original graph in such a way that there should be at least k-1 other candidate nodes with similar features (the features are part of the assumption made on the capability of the attacker).

In both cases one can assure that re-identification cannot be more precise that randomly picking among at least k candidates. It looks natural to apply our algorithms to the clustering approach (actually Hay

et al([9]) suggested a similar approach for social network anonymization see section 6).

5 Experimental Results

Programs have been implemented with C++, OpenMP. The hardware configuration for experiments is given in **Table 1**.

Name	Crunch1
CPU	4 x Intel(R) Xeon(R) CPU E5-4620 0 @ 2.20GHz
Total cores	32
Memory(GB)	379
System type	PowerEdge R820 (Dell Inc.)

Table 1: Hardware configuration

5.1 Comparison between MMSBM and MCMMSBM

The benchmark for the comparison between **MMSBM** and **MCMMSBM** the recommender system (see **Section 4**). The theoretical basis of both **MMSBM** and **MCMMSBM** is similar same. The algorithms are slightly different. Both are using soft clustering, here, we assume that rating between two clusters U_i, and V_j follows *multinomial* distributions. we used the 100k movielens dataset [8]. For the evaluation of the predictions, 5-fold cross validation are done. First, let's compare how these algorithms optimize entropy. Results are depicted in **Figure 3** (results given for 10 clusters for both of user and product clusters).

As you can see in the **Figure 3**, **MCMMSBM** optimizes entropy more efficiently than **MMSBM**. Moreover, we can also set a trade-off: if we choose small value for s, for instance $s(=10)$, then **MCMMSBM** converges in the fastest way but overall it converges too fast. For $s = 30$, entropy is even smaller than **MMSBM** at time 300sec. Thus, one can *flexibly* choose s by considering hardware environment or time available. On the other hand, if

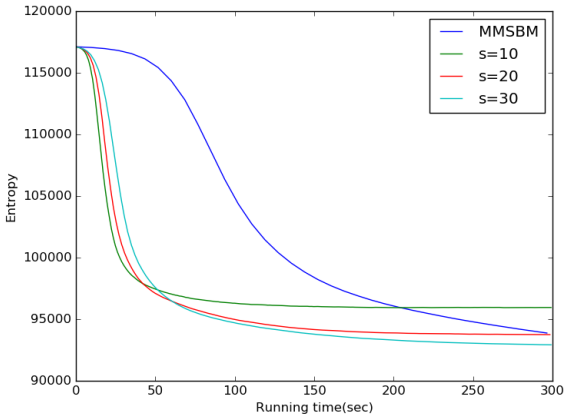

Figure 3:

data size becomes huge, the time cost of one iteration is very expensive in **MMSBM**. Moreover, if we choose bigger k, l, then **MCMMSBM** becomes more efficient than **MMSBM**.

Now, let's evaluate the prediction. The measurment used is the RMSE (Root Mean Square Error). For given pair of user and product in test set, we can estimate the rating as **Eq (13)**.

$$\tilde{r}(u,v) = \sum_{r} \sum_{i,j} h_{u,i} \cdot h_{v,j} \cdot \theta_{i,j,r} \cdot r \qquad (13)$$

Then, the RMSE is defined as **Eq (14)**, where Y is the test set.

$$RMSE = \sqrt{\frac{1}{|Y|} \sum_{(u,v)\in Y} |r(u,v) - \tilde{r}(u,v)|^2} \qquad (14)$$

The result of RMSE is depicted as **Table 2**, we can see that small entropy implies nice prediction.

5.2 Comparison between soft, and hard clustering

In order to compare soft, and hard clustering, we also applied hard clustering to recommender system. Here each user, and product belongs to one of cluster *deteministically*. The degree of freedom is smaller

	Entropy	RMSE
MMSBM	93876.4	0.9536
$s = 10$	95952.2	0.9584
$s = 20$	93748.6	0.9515
$s = 30$	92920.0	0.9510

Table 2: Running time: 300sec

than for soft clustering, thus the optimized entropy of training set will be bigger. But, optimizing entropy is NOT our goal, don't forget that our goal is *prediction*. We can still expect better prediction with hard clustering, despite of bigger entropy.

First, let's compare the optimization of entropy. Results are depicted in **Figure 4**. The number of clusters is set to 15 both for user and product clusters. We can see that hard clustering (Generalized k-means) is much faster than soft clustering, and that optimized entropy is bigger, as it was expected.

Next, let's evaluate prediction using RMSE. Results are depicted in **Table 3**. The performance of prediction is also worse than soft clustering. In summary: hard clustering is cheaper, with a wrost quality than soft clustering.

	Entropy	RMSE	Runni
Generalized k-means	97741.0	0.9713	2
MCMMSBM ($s = 30$)	92920.0	0.9510	3(

Table 3:

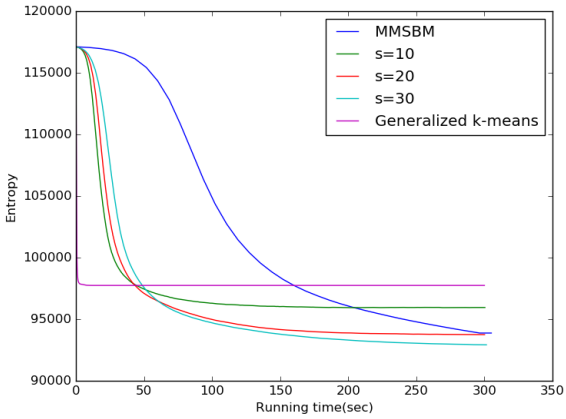

(a)

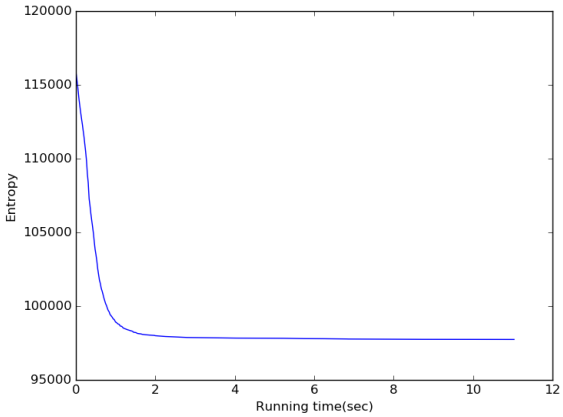

(b)

Figure 4:

5.3 Result of social network anonymization

Applying soft clustering to social network anonymization may be untractable, because even for regenerating random network, it has a complexity in $O(n^2)$, while in hard clustering the complexitiy is in $O(m)$, where n is the number of vertices, and m is the number of edges (*cf.* the complexity of generating ER random graph is $O(n \cdot p)$). Thus here, we only apply hard graph clustering to social network anonymization with Generalized k-means algorithm.

We used two networks from KONECT [12] as benchmark of this subsection. First network is Caida network. This is the undirected network of autonomous systems of the Internet connected with each other from the CAIDA project, collected in 2007. Nodes are autonomous systems (AS), and edges denote communication ($|V(G)| = 26,475, |E(G)| = 53,381$).

The second network is the arXiv astro-ph network. This is the collaboration graph of authors of scientific papers from the arXiv's Astrophysics (astro-ph) section. An edge between two authors represents

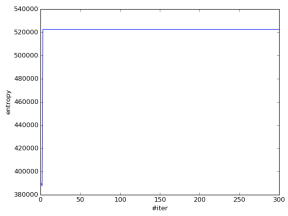

(a) $\alpha = 1.0$

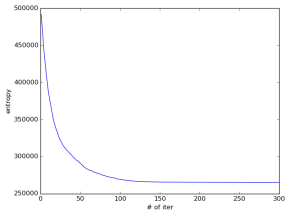

(b) $\alpha = 0.1$

Figure 5: Caida network, 100-clustering, #iter=300, Running time=10.6 s

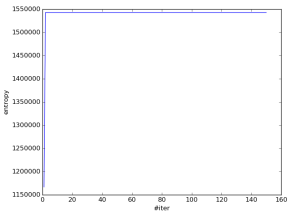

(a) $\alpha = 1.0$

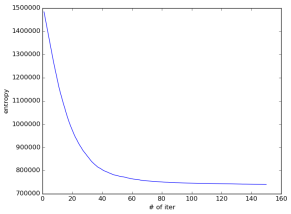

(b) $\alpha = 0.1$

Figure 6: arXiv astro-ph network, 300-clustering, #iter=150, Running time=26 s

a common publication ($|V(G)| = 18,771, |E(G)| = 198,050$).

First, we show that using big iteration scale α doesn't work in order to optimize entropy. Let us consider **Figure 5** and **Figure 6**. We can see that using $\alpha = 1.0$ doesn't work, as we argued previously, the distortion during the parallel loop becomes an issue if we take a large α. On the other hand, when a small α is prefered (0.1), the distortion is negligible, so that our algorithm optimizes entropy very well and fast.

Next, we evaluate the similarity between the original network and the anonymized network. There are various measurements for evaluating network similarity, here, we will compare APL (Average Path Length), GCC (Global Clustering Coefficient), Degree distribution.

The similarity is compared by changing the number of clusters k. In **Figure 7** are plotted the entropy of optimized clustering and random clustering. We can see the number of clusters doesn't affect much to entropy in random clustering, while in optimized clustering, a larger number of clusters implies a better entropy.

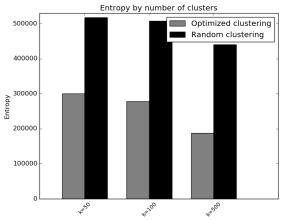

(a) Caida

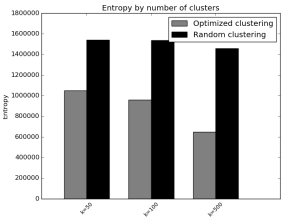

(b) arXiv-AstroPh

Figure 7: Entropy

Now let's see how this entropy affect the general

properties of the graphs. Since anonymized network are *randomly* generated, we have generated them 5 times, and then estimated APL, GCC for each network. Thus the error bar is also represented in **Figure 8, 9**, but we can see that such error bars are very tight.

6 Related Works

MCMMSBM (Algorithm 3) is closely related to the **MMSBM** (originally defined in [1]) and more precisely to the algorithm of [7]: see **Algorithm 8**.

Algorithm 8: Mixed Memebership SBM(MMSBM)

Input : observed edge data X, $V(G)$: set of vertices, k: number of clusters
Output: optimal θ, h

1 Take random $\hat{\theta}, \hat{h}$
2 **repeat**
3 **for** $(u,v) \in X$, *and* $i,j \in [k]$ **do in parallel**
4 $x_{u,v}(i,j) \leftarrow \frac{h_{u,i} \cdot h_{v,j} \cdot \theta_{i,j,r(u,v)}}{\sum_{i',j' \in K} h_{u,i'} \cdot h_{v,j'} \cdot \theta_{i',j',r(u,v)}}$
5 **end**
6 **for** $u \in V(G), i \in [k]$ **do in parallel**
7 $h_{u,i} \leftarrow \frac{1}{|\partial u|} \cdot \sum_{v \in \partial u} \sum_{j \in K} x_{u,v}(i,j)$
8 **end**
9 **for** $i,j \in [k], r \in R$ **do in parallel**
10 $\theta_{i,j,r} \leftarrow \frac{\sum_{(u,v) \in X, r(u,v)=r} x_{u,v}(i,j)}{\sum_{(u,v) \in X} x_{u,v}(i,j)}$
11 **end**
12 **until** $S(\theta, h)$ *converges*;

Note that lines 6∼8 correspond to the M-step of EM algorithm, instead of the E-step. Actually, the lines 3∼5 correspond to the E-step, and lines 6∼11 correspond to the M-step. The time complexity of algorithm 8 is $O(|X| \cdot k^2)$. If k is considered as constant, the complexity is thus $O(|X|)$. However, if the size of data become huge, then the number of clusters k may also increase. Moreover, k^2 is *never* negligible, even if k is small.

On the other hand, the time complexity of **MCMMSBM** is $O(|X| \cdot (k + s \cdot log(k)))$, the sample size s being chosen by trade-off between cost and accuracy.

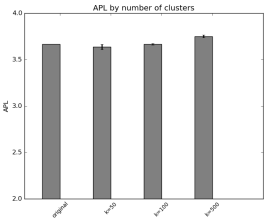

(a) Caida, optimized clustering

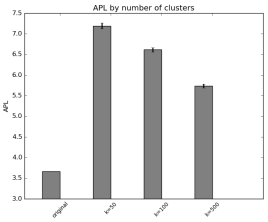

(b) Caida, random clustering

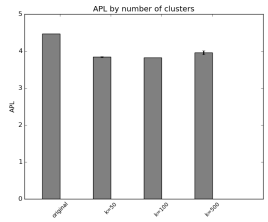

(c) arXiv-AstroPh, optimized clustering

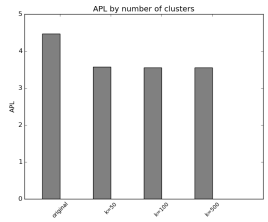

(d) arXiv-AstroPh, random clustering

Figure 8: APL

The update of $h_{u,i}$ in **Algorithm 8** (line 7) is computed as follows:

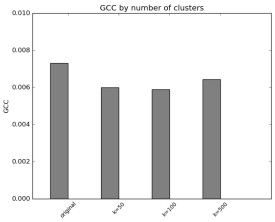

(a) Caida, optimized clustering

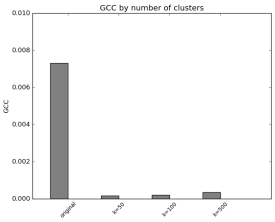

(b) Caida, random clustering

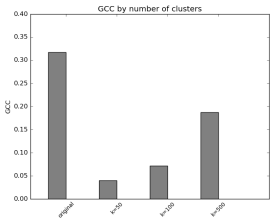

(c) arXiv-AstroPh, optimized clustering

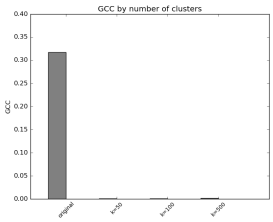

(d) arXiv-AstroPh, random clustering

Figure 9: GCC

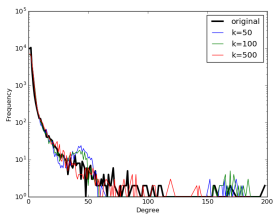

(a) Caida, optimized clustering

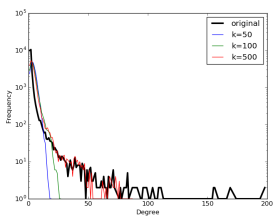

(b) Caida, random clustering

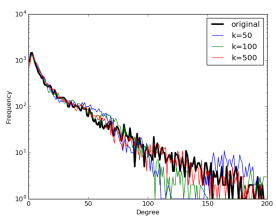

(c) arXiv-AstroPh, optimized clustering

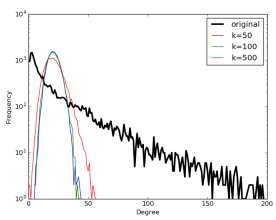

(d) arXiv-AstroPh, random clustering

Figure 10: Degree distribution

$$
\begin{aligned}
h_{u,i} =& \frac{1}{|\partial u|} \cdot \sum_{v \in \partial u} \sum_{j \in K} x_{u,v}(i,j) \\
=& \frac{1}{|\partial u|} \cdot \sum_{v \in \partial u} \frac{\sum\limits_{j' \in K} h_{u,i} \cdot h_{v,j'} \cdot \theta_{i,j',r(u,v)}}{\sum\limits_{i',j' \in K} h_{u,i'} \cdot h_{v,j'} \cdot \theta_{i',j',r(u,v)}}
\end{aligned} \tag{15}
$$

The computation of the denominator requires a computation in $O(k^2)$.

If $h_{u,i'} \cdot h_{v,j'}$ is considered as a pmf of some random variable we have to define the following random

variables:

$$
\begin{aligned}
X_{u,v} &:= \theta_{i',j',r(u,v)} \text{ w/p } h_{u,i'} \cdot h_{v,j'} \\
X_{u,v}(i) &:= \begin{cases} \theta_{i,j',r(u,v)} & \text{w/p } h_{u,i} \cdot h_{v,j'} \\ 0 & \text{w/p } 1 - h_{u,i} \end{cases} \\
X_{u,v}(i,j) &:= \begin{cases} \theta_{i,j,r(u,v)} & \text{w/p } h_{u,i} \cdot h_{v,j} \\ 0 & \text{w/p } 1 - h_{u,i} \cdot h_{v,j} \end{cases}
\end{aligned} \tag{16}
$$

where w/p means "with probability", and $i', j' \in [k]$, then

$$
\begin{aligned}
\sum_{i',j' \in [k]} h_{u,i'} \cdot h_{v,j'} \cdot \theta_{i',j',r(u,v)} &= \mathbb{E}[X_{u,v}] \\
\sum_{j' \in [k]} h_{u,i} \cdot h_{v,j'} \cdot \theta_{i,j',r(u,v)} &= \mathbb{E}[X_{u,v}(i)]
\end{aligned} \tag{17}
$$

$X_{u,v}, X_{u,v}(i), X_{u,v}(i,j)$ can be generated as follows:

1. Pick i', j' respectively from $discrete(h_u), discrete(h_v)$
2. Assign the value of random variables $X_{u,v}, X_{u,v}(i), X_{u,v}(i,j)$ as:

$$
\begin{aligned}
X_{u,v} &:= \theta_{i',j',r(u,v)} \\
X_{u,v}(i) &:= \begin{cases} \theta_{i',j',r(u,v)} & \text{if } i' = i \\ 0 & \text{else} \end{cases} \\
X_{u,v}(i,j) &:= \begin{cases} \theta_{i',j',r(u,v)} & \text{if } i' = i \text{ and } j' = j \\ 0 & \text{else} \end{cases}
\end{aligned} \tag{18}
$$

Those random variables can be sampled in $O(k + s \cdot log(k))$ for sample size s. Indeed, it is a joint distribution of *independent* events. For $s' \in [s]$, we can sample

$$
\begin{aligned}
X_{u,v,s'} &\overset{i.i.d}{\sim} X_{u,v} \\
X_{u,v,s'}(i) &\overset{i.i.d}{\sim} X_{u,v}(i) \\
X_{u,v,s'}(i,j) &\overset{i.i.d}{\sim} X_{u,v}(i,j)
\end{aligned} \tag{19}
$$

Thus through Monte-Carlo simulation, g, θ can be approximately updated as follows:

$$
\begin{aligned}
h_{u,i} &= \frac{1}{|\partial u|} \sum_{v \in \partial u} \frac{\mathbb{E}[X_{u,v}(i)]}{\mathbb{E}[X_{u,v}]} \\
&\approx \frac{1}{|\partial u|} \sum_{v \in \partial u} \frac{\frac{1}{s}\sum_{s'} X_{u,v,s'}(i)}{\sum_{s'} X_{u,v,s'}} \\
\theta_{i,j,r} &= \frac{\sum_{(u,v) \in X, r(u,v) = r} x_{u,v}(i,j)}{\sum_{(u,v) \in X} x_{u,v}(i,j)} \\
&\approx \frac{\sum_{(u,v) \in X, r(u,v) = r} \tilde{x}_{u,v}(i,j)}{\sum_{(u,v) \in X} \tilde{x}_{u,v}(i,j)}
\end{aligned} \tag{20}
$$

where

$$
\tilde{x}_{u,v}(i,j) := \frac{\frac{1}{s}\sum_{s'} X_{u,v,s'}(i,j)}{\frac{1}{s}\sum_{s'} X_{u,v,s'}} \tag{21}
$$

It is computed in $O(|X| \cdot (k + s \cdot log(k)))$ time by **Algorithm 3**.

The application of the Generalized k-means algorithm to the problem of social network anonymization amounts to an approach similar to the one of proposed in [9] by Hay *et al.* We can note the following differences between two approaches are:

1. In [9] the edges between two clusters are just rearranged, thus the number of edges between two clusters is a constant. With our approach the number may change.
2. With our approach, a totally new random graph is generated which is only *similar* to the original one. The concept of k-anonymity cannot be strictly applied. Namely we don't have to fix the minimum size of clusters which differs with [9].

Moreover, the experimental results shows a very large difference in terms of efficiency between our approach and [9]. Despite different hardware configuration with [9], Generalized k-means appears much faster.

For example in [9] the following result is given: it takes 1 hour to cluster graphs of size 5000, on the other hand it takes 10 seconds using Generalized k-means to cluster graphs of size 30000.

7 Conclusion

In this paper we have presented the **MCMMSBM** algorithm which can be seen as an improved **MMSBM** algorithm [7] by applying Monte-Carlo simulation to point of *efficiency*. Theoretically, **MCMMSBM** can't achieve better optimizations than **MMSBM**, if one considers *infinite* computing resources. But in reality, **MMSBM** is *strictly* limited. We have shown that **MCMMSBM** can achieve better optimization, and better prediction in **Section 5.1**.

We also proposed the Generalized k-means algorithm. It can be widely applied for hard classification problems, especially for hard graph clustering problems. We reclassified small proportion of data (or nodes), instead of the whole data in one iterration.

We have applied SBM to social network anonymization. We saw that entropy optimization works very well for property preservation. We also compared the results by changing the number of clusters k, we can consider that there can be trade-off for deciding k. If we take small k, we saw that network properties are not preserved well. But if we take large k, *anonymity* can be vulnerable. We left as future work deeper comparisons with other social network anonymization techniques, notably on the quality of the published network, for instance with [15, 4].

In order to measure the anonymity achieved, we have relied on the k-anonymity [2] definition which is widely applied. But such definition is not really suitted for this anonymization framework. As future work we consider to define an appropriate anonymity measurement for this framework.

References

[1] E. M. Airoldi, D. M. Blei, S. E. Fienberg, and E. P. Xing. Mixed membership stochastic blockmodels. *Journal of Machine Learning Research*, 9:1981–2014, 2008.

[2] L. Backstrom, C. Dwork, and J. M. Kleinberg. Wherefore art thou r3579x?: anonymized social networks, hidden patterns, and structural steganography. *Commun. ACM*, 54(12):133–141, 2011.

[3] J. Bobadilla, F. Ortega, A. Hernando, and A. Gutirrez. Recommender systems survey. *Knowledge-Based Systems*, 46:109 – 132, 2013.

[4] A. Campan, Y. Alufaisan, and T. M. Truta. Preserving communities in anonymized social networks. *Trans. Data Privacy*, 8(1):55–87, 2015.

[5] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *JOURNAL OF THE ROYAL STATISTICAL SOCIETY, SERIES B*, 39(1):1–38, 1977.

[6] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *JOURNAL OF THE ROYAL STATISTICAL SOCIETY, SERIES B*, 39(1):1–38, 1977.

[7] A. Godoy-Lorite, R. Guimerà, C. Moore, and M. Sales-Pardo. Accurate and scalable social recommendation using mixed-membership stochastic block models. *CoRR*, abs/1604.01170, 2016.

[8] F. M. Harper and J. A. Konstan. The movielens datasets: History and context. *TiiS*, 5(4):19, 2016.

[9] M. Hay, G. Miklau, D. Jensen, D. F. Towsley, and C. Li. Resisting structural re-identification in anonymized social networks. *VLDB J.*, 19(6):797–823, 2010.

[10] P. W. Holland, K. B. Laskey, and S. Leinhardt. Stochastic blockmodels: First steps. *Social Networks*, 5(2):109 – 137, 1983.

[11] J. Kleinberg and E. Tardos. Approximation algorithms for classification problems with pairwise relationships: Metric labeling and markov random fields. *J. ACM*, 49(5):616–639, Sept. 2002.

[12] J. Kunegis. KONECT: the koblenz network collection. In *22nd International World Wide Web Conference, WWW '13, Rio de Janeiro, Brazil, May 13-17, 2013, Companion Volume*, pages 1343–1350, 2013.

[13] A. Narayanan and V. Shmatikov. De-anonymizing social networks. In *30th IEEE Symposium on Security and Privacy (S&P 2009), 17-20 May 2009, Oakland, California, USA*, pages 173–187, 2009.

[14] S. E. Schaeffer. Graph clustering. *Computer Science Review*, 1(1):27–64, 2007.

[15] T. M. Truta, A. Campan, and A. L. Ralescu. Preservation of structural properties in anonymized social networks. In C. Pu, J. Joshi, and S. Nepal, editors, *8th International Conference on Collaborative Computing: Networking, Applications and Worksharing, CollaborateCom 2012, Pittsburgh, PA, USA, October 14-17, 2012*, pages 619–627. ICST / IEEE, 2012.

[16] J. J. Verbeek, N. Vlassis, and B. Krse. Efficient greedy learning of gaussian mixture models. *Neural Computation*, 15:469–485, 2003.

[17] C. F. J. Wu. On the convergence properties of the em algorithm. *Ann. Statist.*, 11(1):95–103, 03 1983.